

Feedback interactions between neuronal pointers and maps for attentional processing

Richard Hahnloser^{1,3}, Rodney J. Douglas¹, Misha Mahowald^{1*} and Klaus Hepp²

¹ Institute of Neuroinformatics, ETHZ/UNIZ, CH-8057 Zürich, Switzerland

² Institute for Theoretical Physics, ETHZ, CH-8093 Zürich, Switzerland

³ Department of Brain and Cognitive Science, MIT E25-210, Cambridge, Massachusetts 02139, USA

Correspondence should be addressed to R.H. (rh@ai.mit.edu).

Neural networks combining local excitatory feedback with recurrent inhibition are valuable models of neocortical processing. However, incorporating the attentional modulation observed in cortical neurons is problematic. We propose a simple architecture for attentional processing. Our network consists of two reciprocally connected populations of excitatory neurons; a large population (the map) processes a feedforward sensory input, and a small population (the pointer) modulates location and intensity of this processing in an attentional manner dependent on a control input to the pointer. This pointer-map network has rich dynamics despite its simple architecture and explains general computational features related to attention/intention observed in neocortex, making it interesting both theoretically and experimentally.

The fundamental similarities in structural organization and physiology across neocortex suggest that the cortical circuits use common principles of operation that can be modified according to specific processing tasks^{1,2}. One fundamental principle is the feedback excitation^{3–5} mediated by recurrent intracortical axonal connections^{2,6–8}.

The functional properties of recurrence are often studied in network models consisting of a one-dimensional (1-D) array of neurons, each with a similar pattern of excitatory and inhibitory connections^{2,9–12}. Typically, the location in the array of responding neurons encodes a stimulus variable (for instance, spatial location or orientation). We will call such a network a recurrent ‘map’ (Fig. 1a).

Recurrent maps have interesting properties that contribute directly to our understanding of signal processing by cortical neuronal networks. In particular, excitatory feedback generated between map neurons enhances those features of the input matching the synaptic patterns of recurrence, and overall inhibition suppresses noise in these patterns¹³. However, the recurrent map is unable to focus processing on some specific aspect or region of its input, and so fails to incorporate the attentional/intentional modulation observed in many cortical networks^{14–18}.

We incorporated the required attentional mechanism in the recurrent map by inserting a small population of ‘pointer’ neurons in the excitatory feedback loop of the map (compare Fig. 1a with b). In this case, the recurrent excitation on the map depends not only on the synaptic weights, but also on the activation of the pointer neurons. Consequently, top-down inputs to the pointer neurons that affect their activation can control the map response to sensory input in an attentional manner by modulating excitatory feedback. We call this new recurrent architecture a ‘pointer map’ (Fig. 1b).

*deceased

RESULTS

Our results are of two kinds, mathematical (see Methods and Mathematical Appendix) and simulation. The mathematical results address some general properties of the pointer-map dynamics such as their convergence (to fixed point attractors) and their relationship to the dynamics of recurrent maps.

Our simulation results are illustrated for a pointer map encoding a stimulus feature, for example, the retinal location of a visual target. The activities of the 25 neurons in the one-dimensional map are labeled by $\mathbf{M} = (M_1, \dots, M_N)$. Usually, the sensory input $\mathbf{m} = (m_1, \dots, m_N)$ to the map consisted of either one or two stationary or slowly moving visual stimuli. A visual stimulus at location r was modeled by a Gaussian activity profile $m_x = h \exp[-(r - \delta_x)^2 / \sigma^2]$, where the height h and width σ reflect time-invariant properties of the stimulus (for example, size, contrast, and orientation). Here, $0 \leq \delta_x \leq \pi/2$ is the preferred spatial stimulus location in retinal (angular) coordinates of the x th map neuron.

When the visual stimulus was stationary, the activities $\mathbf{P} = (P_1, P_2)$ of the two pointer neurons reported its presence by virtue of their sine and cosine synaptic weightings with the map:

$$\mathbf{P} = (P_1, P_2) = [\alpha \sum_x (\cos \delta_x, \sin \delta_x) M_x + (p_1, p_2)]_+ \quad (1)$$

Here $[\cdot]_+ = \max(0, \cdot)$ denotes a threshold nonlinearity that makes pointer activities nonnegative. We included a parameter α in our model that uniformly scaled the strength of individual excitatory connections between the map and pointer neurons as well as a similar parameter β for the global inhibition on the map to explore the effects of weak and strong synaptic connections.

We denote an attentional input to the pointer neurons by $\mathbf{p} = (p_1, p_2)$. For simplicity of exposition, we first assume that this attentional input is zero. In this case, because of our particular choice of synaptic weighting, the activities of the pointer neu-

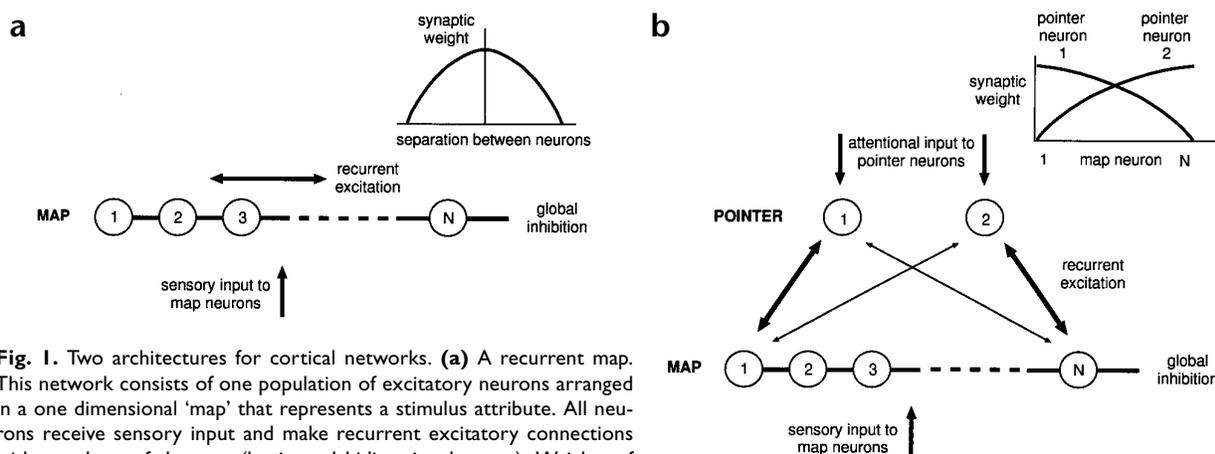


Fig. 1. Two architectures for cortical networks. **(a)** A recurrent map. This network consists of one population of excitatory neurons arranged in a one dimensional 'map' that represents a stimulus attribute. All neurons receive sensory input and make recurrent excitatory connections with members of the map (horizontal bidirectional arrow). Weights of synaptic connections fall off with distance between neurons (inset). All neurons receive the same global inhibitory input (solid line drawn through map neurons), which is proportional to the summed activity on the map. **(b)** A pointer map. Like the recurrent map, this network comprises a population of excitatory neurons. However, unlike the recurrent map, recurrent excitatory connections do not occur directly between members of the map, but through a second population consisting of two of excitatory 'pointer' neurons. The two pointer neurons make different connection patterns with the map neurons. Pointer neuron 1 makes identical forward and backward connections that are a cosine function of location on the map (inset), giving stronger connections on the left (thick bidirectional arrow) than on the right (thin bidirectional arrow). Pointer neuron 2 makes connections that are a sine function of location, and so increase toward the right. The pointer neurons also receive an attentional input.

rons form a population vector of steady map activity. The angle $\gamma = \arctan(P_2/P_1)$ formed by the pointer reports the location of the visual stimulus. The length $\sqrt{P_1^2 + P_2^2}$ of the pointer reports the intensity of the stimulus, or possibly some other variable (such as eye position) encoded by the network.

A vector representation of the activities of a population of neurons is useful for decoding the firing rates of cortical neurons^{19,20}. However, such a population vector is a passive read-out of activity on the map. By contrast, because of recurrence, our pointer contributes actively to processing on the map. That is, the pointer neurons excite neurons on the map whose preferred stimulus locations are closest to the direction γ of the pointer (Fig. 2). One of the effects of the excitatory and inhibitory feedback in the network is to localize the activity on the map by suppressing lateral noise. Most importantly, the attentional input p to the pointer neurons biases the pointer in equation 1. We propose processing induced by this attentional input as an alternative architecture that can be combined with recurrent maps. The attentional input imposes top-down selectivity on processing by the network, affecting the pointer and, indirectly, the map activity. As a consequence, sensory responses of both pointer and map neurons reflected attentional modulation, but at different spatial scales. The map neurons had a relatively narrow spatial tuning, whereas the tuning of the pointer neurons was broader.

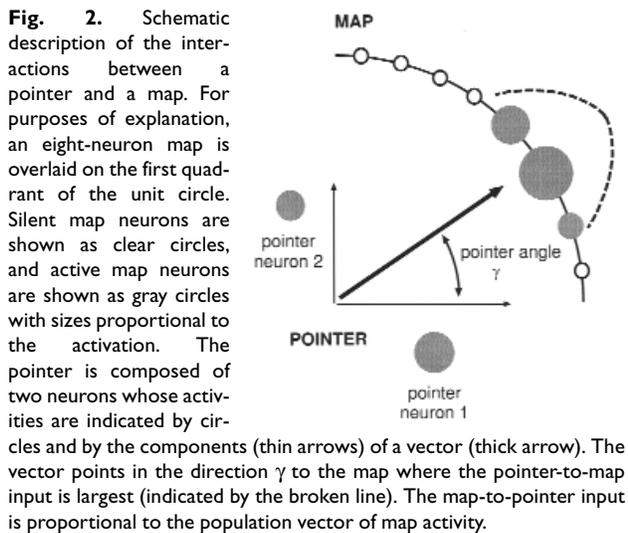
Figure 3a shows the steady response of a pointer map to a degraded static visual image. The combination of excitatory and inhibitory feedback enhances the signal, suppressing noise in the input. The stronger the feedback (the larger α and β), the sharper the stationary activity profile. For both weak and strong feedback, the pointer is directed toward the spatial location of the stimulus.

When two or more visual stimuli were presented to the network, strong feedback could restrict the representation on the map to just one of them (Fig. 3b). The representations of spatially distant stimuli compete with one another, and the network can be multistable. That is, the most salient stimulus did not nec-

essarily win the competition. The stronger the feedback, the more likely that a weak stimulus could be stably represented. The winning stimulus depends on the initial activations of map and pointer neurons. In particular, the preattentive bias of the competition induced by the initial activation of pointer neurons can be used to initialize the processing of visual input to an expected, specific, spatial location of a visual stimulus. In addition, for weak feedback (not shown), several stimuli could be simultaneously represented by the map and the pointer was directed toward their average location.

The attentional input p to the pointer can change the location and the magnitude of the map activity profile even after presentation of a visual stimulus. The effects of the attentional input can be decomposed into two orthogonal components relative to the instantaneous direction of the pointer P . The component of the input p orthogonal to P has a directing or steering effect, causing rotation of the pointer. Figure 4a shows how the control input could be used to direct the pointer from the right toward the left, forcing the selection of the weaker of two visual stimuli. Because the pointer map is multistable, the weaker input can remain selected even after the control input has been removed. In this example, we applied an excitatory input to the left pointer neuron. A similar steering could be achieved by applying an inhibitory input to the right pointer neuron. Although these two methods produced the same result of selecting the stimulus to the left, they are dynamically different. An excitatory input to the left pointer neuron tends to activate map neurons whose preferred stimulus locations are on the way to the stimulus to be selected, whereas a negative input to the right pointer neuron tends to inhibit them.

Figure 4b demonstrates movement of an activity profile about the map by the pointer. The map was driven by uniform input, arising from non-specific visual input and/or from an efference copy of eye position (see below). The network transformed the uniform input into a narrow profile of activity on the map. The activity could then be shifted around continuously by steering



the pointer. This is a simple way of implementing dynamic remapping²¹, the internally generated shifting of activity profiles, as observed in premotor and parietal cortex^{22,23}, for example.

Besides rotating the pointer, the attentional input can either facilitate or suppress. The component of \mathbf{p} parallel to the pointer causes facilitation of the pointed-to region on the map. If the component is antiparallel (inhibitory), then the corresponding region is suppressed. The facilitatory input to the pointer could be used, for example, to hold activity at an attended/intended location on the map (Fig. 4c). In this example, the visual input contained a target that first drifted toward the left, held position and then drifted toward the right again. Initially, the pointer simply tracked the target location. When the target was on the left, the pointer received a facilitatory attentional input causing the map activity to remain latched even after the target had drifted away. Strong inhibitory attentional input to pointer neurons could be used to derecruit excitatory feedback. In Fig. 4d, we connected two pointer neurons to a recurrent map. The strong inhibitory input ($p_1, p_2 < 0$) completely suppressed pointer neurons and caused only the recurrence to be effective and represent the stronger stimulus. Alternatively, by disinhibiting the pointer neu-

rons (no attentional input), there was an increased positive feedback on the map that caused higher map activity. Although the pointer neurons were only weakly connected to the map in this example, target selection caused by attentional input as in Fig. 4a could still be achieved.

Pointer-map interactions offer a simple explanation of how a neuron can be both locally tuned to some stimulus variable and monotonically tuned to some other variable. This is commonly observed in neurons of parietal areas LIP and 7a, which combine a localized retinal receptive field multiplicatively with a monotonic tuning to eye position^{24,25}, for example. In the pointer map, such multiplicative response properties can arise in two ways (Fig. 5). One mechanism (Fig. 5a and b) acted using map input as in a recurrent map^{11,26}. The other mechanism (Fig. 5c and d) acted using the pointer input. In the case of map input, addition of a uniform feedforward input e (an efference copy of eye position) to a Gaussian visual input was expressed as a multiplicative combination in the output. The reason for this transformation is the localizing effect of excitatory feedback. The Gaussian (visual/saccadic) input establishes the output response location, and the uniform input scales the localized response. Because the pointer and map are linked by feedback, the modulation of response by uniform input affects both map and pointer neurons, although the modulatory input is applied to the map neurons only.

In the pointer map, attentional input to the pointer could also mediate multiplicative scaling (Fig. 5c and d). In this case of Fig. 4c, however, the facilitatory attentional input had to be parallel to the pointer rather than uniform. The attentional input intensity scaled the activity profile on the map. For a particular map neuron, this scaling is expressed as an attentional modulation of the gain of its visual receptive field. The two methods of multiplication, homogeneous input e to the map and attentional input to the pointer, produced almost indistinguishable gain modulations (compare Fig. 5a to c). By combining these two mechanisms, the firing rates of the neurons can be influenced independently by two inputs. For example, if the map receives several visual targets, then the focal pointer input can address and selectively enhance one of these targets while suppressing others. The constant additive eye position input to the map scales the amplitude of the winning target in a multiplicative fashion. This mechanism solves the conundrum¹⁶ of how an ensemble of retinocentric parietal neurons

Fig. 3. Pointer map without attentional input ($\mathbf{p} = \mathbf{0}$) (a) Signal restoration by pointer-map feedback. The noisy visual input \mathbf{m} to the map is shown by the dashed line. It consists of a Gaussian centered on neuron 11 to which a uniform Gaussian noise with mean 0 and variance 0.5 was added. The steady map activity is shown for weaker ($\beta = 0.1$, $\alpha = 0.34$, thin line) and stronger ($\beta = 10$, $\alpha = 3.16$, thick line) feedback connections, resulting in weaker and stronger competition between distant neurons on the map. In both cases, the steady pointer (shown by the arrow) is directed toward the location of the visual input on the map, correctly extracting the location of the noisy visual stimulus. (b) Competing dynamics. Shown is the input (below) and time evolution of activity (above) of the 25 map neurons (left) and 2 pointer neurons (right). All activities are represented in pseudocolor with saturated blue at zero increasing to saturated red. The location on the map corresponding to the direction $\gamma(t)$ of the pointer is shown as a white line superimposed on the map output. The map input contains a weak target on the left and a stronger target that wins the competition on the right. The map and pointer neurons are initialized to zero activity. $h_1 = 0.5$, $h_2 = 0.5$, $\sigma = \sqrt{5}$ and $\beta = 1$.

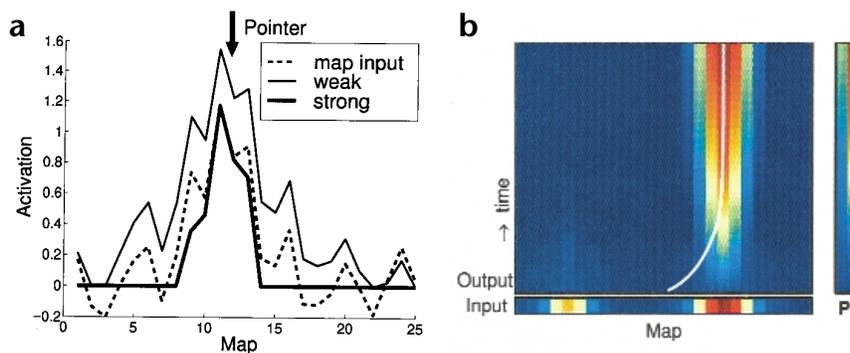
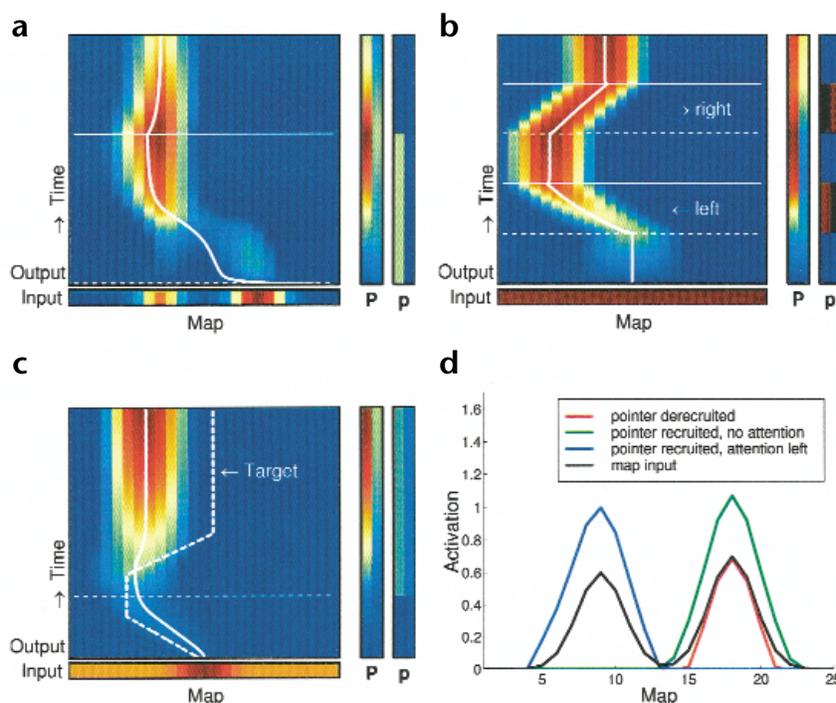


Fig. 4. Attentional pointer map. The time-dependent activities of the two pointer neurons, \mathbf{P} , and their inputs, \mathbf{p} , are shown at the right. Dashed horizontal lines mark the time when attention is engaged, and full horizontal lines mark the time when attention is withdrawn ($\mathbf{p} = \mathbf{0}$). $h_1 = 0.7$, $h_2 = 0.5$ and $\sigma = \sqrt{5}$. **(a)** Target selection. The pointer is initialized to the right, $\mathbf{P}(0) = (0, 1.5)$. Thereafter, the attentional input $\mathbf{p} = (0.6, 0)$ forces the pointer to direct its excitatory feedback loop to the left of the map, bypassing the location of strongest input and selecting the weaker input to the map. Even after withdrawal of attention, the activity profile remains at the weaker input, with just a small shift to the right. $\beta = 3$. **(b)** Dynamic remapping. The input to the map is uniform. Initially, there is no external input to the pointer, so the attractor in the center of the map is selected. During the period marked 'left', the pointer receives an input $\mathbf{p} = (0.4, -0.4)$, which steers the activity on the map continuously toward the left. During the period marked 'right', the opposite input to the pointer $\mathbf{p} = (-0.4, 0.4)$ steers the activity toward the right. When the pointer input is withdrawn, the activity settles down at its current map location. The inhibitory input to pointer neurons is shown in black. $\beta = 3$. **(c)** Attentional memorization.

The map receives a uniform background input of amplitude 1 that provides the energy that enables self-sustained activity. A Gaussian target of height 1 is added to the input. The location of the target evolves in time (dashed line). First, it drifts from the middle toward the left, where it stops. The passive pointer tracks the target. At a certain time t_0 , the pointer receives a sustained input parallel to its current direction, $\mathbf{p} = \frac{1}{5} \mathbf{P}(t_0)$. When the target subsequently starts to drift away toward the right, the map activity and the pointer remain at the attended location. $\beta = 4$. **(d)** Recruitment of pointer neurons. This figure shows simulation results of a map containing recurrent excitatory map connections with a Gaussian profile, peak synaptic strength of 0.9 and width $\sigma = 10$. Additionally, two pointer neurons are recurrently connected to the map with a synaptic strength of $\alpha = 0.3$. There is global inhibition on the map of strength $\beta = 0.9$. The fixed map input used to compute steady states for three different attentional inputs is shown by the black line. When the pointer neurons are suppressed by inhibitory attentional input $\mathbf{p} = (-0.8, -0.8)$, the response gain of the map is -1 (red line). Only the stronger input is represented by the recurrent map. When the pointer is disinhibited ($\mathbf{p} = \mathbf{0}$) to the pointer, then map gain increases as a result of recruited feedback by the pointer. Here the stronger target leads to higher activity (green line). Finally, attention to the left [$\mathbf{p} = (0.4, -0.4)$] causes selection of the weaker input to the map (blue line).



can simultaneously express the locus of attention and encode eye position by gain modulation.

DISCUSSION

Foveating eyes and reaching limbs are important for focal sensory processing and motor interaction with the world^{23,27}. The direction of their pointing modulates the responses of neurons in a number of cortical and subcortical areas. These observations suggest that neuronal 'pointers' to distributed activity on a map of neurons may be a general computational property of sensorimotor and attentional/intentional processing. The pointer-map network presented here provides a general computational construct in which neuronal pointers operate on associated neuronal maps.

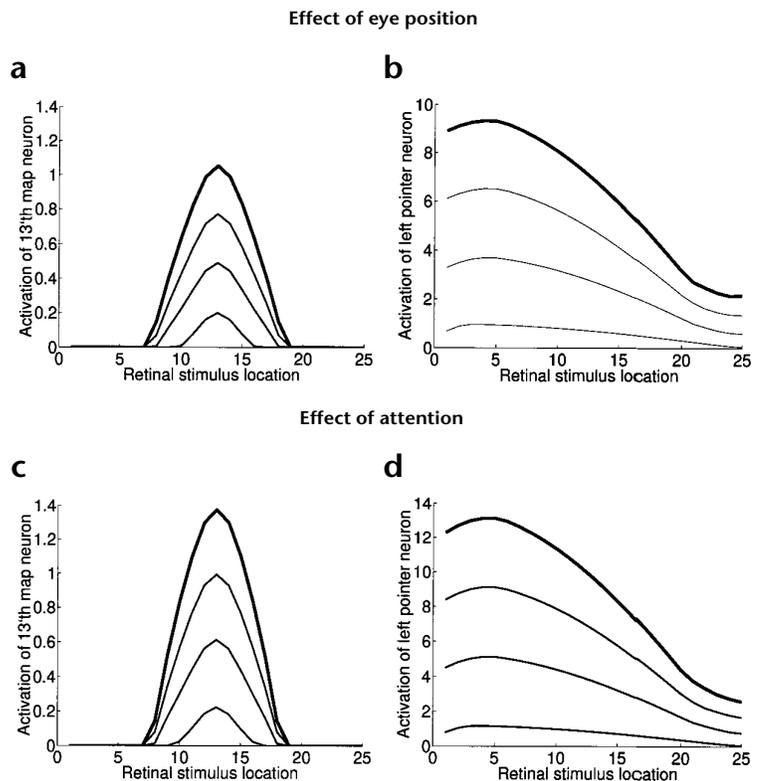
Architecturally, pointer neurons are located in the excitatory feedback loop of the map neurons. The pointer comprises a small number of neurons whose output firing rates forms a vector. When the network is stationary, the direction of the vector points to, and represents, a cluster of active neurons on the map. Attentional inputs to the pointer neurons could steer and modulate the recurrent excitation applied to the map (Figs. 4 and 5). The dynamics of the network depends on interaction of the rectification non-linearity of each neuron with the variable gain arising

from the recurrent connections. The stronger the excitatory and inhibitory feedback was, the stronger the induced competition between different stimuli (Fig. 3a).

There are similarities between pointer maps and recurrent maps. For example, the behavior shown in Fig. 3 has been demonstrated as a feature of recurrent networks^{2,11}. Indeed, in the Mathematical Appendix, we show that if the attentional input \mathbf{p} to the pointer neurons is only positive, then the pointer map can be related to a recurrent map with cosine-tuned synapses that receive additional input, but with a different time dependence. However, the architectural and functional advantages of the pointer map are lost in this reduction because attentional control and read-out are separated from the recurrent map, whereas in the pointer map, both control and readout are incorporated directly into the feedback path of the network itself.

Unlike recurrent maps, our architecture offers the possibility of dynamically changing excitatory feedback strength without synaptic plasticity. By inhibiting pointer neurons, the excitatory feedback to the map could be derecruited (Fig. 4d). As a possible generalization of this nonlinear operation, we can imagine a redundant number (>2) of pointer neurons connected identically to a map. By selectively inhibiting some of them, the excitatory feedback in the network could be continuously reduced,

Fig. 5. Modulation of retinal receptive fields of map and pointer neurons. Steady responses of the 13th map neuron and the left pointer neuron as a function of the retinal location of a visual stimulus. $h = 0.7$, $\sigma = \sqrt{5}$ and $\beta = 2$. **(a, b)** The visual response modulation arises by adding uniform inputs $e = 0, 0.33, 0.66$ and 1 to the map, corresponding to 4 different eye positions, from -45° ($e = 0$) to 45° ($e = 1$). There is no input to the pointer neurons. The largest responses (thick lines) correspond to $e = 1$. **(c, d)** The visual response modulation arises from input $p = p(\cos\theta, \sin\theta)$ to the pointer neurons, where $p = 0.25, 0.5$ and 0.75 correspond to 4 different levels of attentional input. The angle θ is chosen to be identical to the angle γ of the pointer, corresponding to attention at the retinal location of the stimulus. There is a facilitatory effect to the neurons whose preferred stimulus location coincides with the location of the actual visual stimulus. The strongest responses are shown by the thick lines (corresponding to $p = 0.75$). Despite the different mechanisms at work, the map and pointer neurons express similar modulations of the visual response. Because of inhibitory feedback interactions within the map, the receptive fields of map neurons are more narrowly tuned than those of pointer neurons.



leading to stronger or weaker competition between stimuli. This is because the excitatory feedback strength depends only on the number of active pointer neurons.

The pointer-map control can be generalized from maps of one dimension to maps of many dimensions. Topologically, the recurrent weights between the map and at least $n + 1$ pointer neurons are located on the positive quadrant of the n dimensional sphere embedded in $(n + 1)$ -dimensional space. For example, if two spatial dimensions are to be encoded, three pointer neurons can represent the azimuth and the elevation as two independent spatial variables. Important qualitative features such as the selective competition between stimuli are retained. Because the activity of the entire map is encoded in the outputs of just a few pointer neurons, the architecture of the pointer map provides an efficient link between different representational scales, and hierarchical levels. Encoding by pointers reduces the axonal connections required to link inter-area maps from the order $O(N^m)$ to $O(n)$, where N is the number of map neurons required to encode a single dimension. Of course, this optimization of wiring efficiency limits transmitted information to a single location and encoded activity at that location. However, the pointer does permit sequential readout by scanning (Fig. 4a), and so information may be transposed from space to time rather than lost.

The pointer and map neurons responded on very different spatial scales (Fig. 5). This difference in scale electrophysiologically distinguishes pointer from map neurons. In addition, according to our model, pointer neurons have longer response latency to sensory stimuli but shorter latency to attentional effects modulating the sensory responses as compared with map neurons. Interaction of pointers and maps within a common feedback loop provides a natural translation between neuronal encodings of different scales, explaining a principle of intra- and inter-areal feedback connections in cortex. For example, area MT neurons with large receptive fields can be interpreted as pointers that selectively enhance the responses of V1 neurons with smaller fields for regions of visual space. This interpretation is

consistent with experimental results showing that MT-V1 feedback connections modulate V1 responsiveness to stimuli²⁸.

The pointer input p , an attentional control signal, both steered excitatory feedback of the network to a chosen map location (Fig. 4a) and facilitated or suppressed its responses (Fig. 5c). This mechanism can explain changes in sensitivity to stimuli within a single receptive field brought about by selective attention, as observed in visual areas V1, V2, V4, MT and MST^{15,17,29,30}. If map-neuron outputs are collected by a single, hypothetical observer neuron (for instance, in MT), then the neuron's response reports the target attended by the pointer. When pointer direction is changed, the same observer neuron will report the response at a new attended location. Because the pointer is steerable, it can provide the sequential scanning required to resolve ambiguity arising from multiple stimuli³¹.

Previous studies incorporated attentional control to provide a dynamic routing, whereby information is selectively shifted from lower to higher processing levels. In these models, the control signals achieved the routing either by modulating effective connection strength of the feedforward information stream³², or by modulating the gain of the input neurons³³. For example, in one model³³, gains of V4 neurons providing input to an observer neuron in IT are multiplicatively modulated by attention. The pointer map provides a mechanism for this attentional modulation.

Attentional properties of the pointer map are equally relevant in preparation for motor action. Neurons of parietal area LIP remain active when animals withhold their motor response while remembering the location of an extinguished saccade target³⁴. This remembered motor plan can be explained by referring to Fig. 4c, in which neurons of the map encode presaccadic movement fields rather than visual receptive fields. The map response can be sustained even after removal of the stimulus. The multi-stability of pointer maps for constant map input allows both the

map and pointer to latch to their current activities, producing short-term memory.

The pointer map can also explain remapping of responses. For example, the rotation of a population vector during an intentional memorization period as reported in neurons of primary motor cortex²² could be achieved by pointer input (see Fig. 4b). In this case, map neurons would code directions of arm movement.

During fixation, LIP neuron receptive fields are retinocentric. However, their responses are modulated both by the behavioral relevance of the attended stimulus¹⁶ as well as by eye position²⁴. How, then, can a decoding network distinguish between attention and eye position in the response of the population of LIP neurons? This conundrum can be explained by a pointer map of LIP where retinal input carrying many targets induces hills of activity on the map. The pointer control input selects a behaviorally relevant target on the map and suppresses the others. A common additive eye position input to the map induces scaling of the response to the selected stimulus. In this way, the retinocentric coordinates of the visual target and the salient position become bound in the single location of the map response profile, and the amplitude of the response profile reflects eye position. This interpretation is supported by findings showing that the visual world is only sparsely represented in LIP³⁵, with only the most salient or most behaviorally relevant stimulus being strongly represented.

Despite its simple architecture, the pointer map expresses sophisticated computational features of cortical processing such as signal enhancement, sequential readout from maps, attentional control and gainfields. Neural pointers are analog versions of the powerful pointer construct used in computer science³⁶ in the sense that a compact neural pointer is able to select and read out a local structure on a much larger map. Moreover, the pointer map is modular, suggesting that pointer maps can compose more elaborate architectures for cortical computation.

METHODS AND MATHEMATICAL APPENDIX

Pointer-map networks. The pointer-map network (Fig. 1b) comprises simple linear threshold neurons that express the rate coding aspects of cortical neuronal circuits. Outputs of these neurons do not saturate in their range of operation, and so the stability of the network, like cortical networks, depends on feedback inhibition². To simplify the model, recurrent inhibition proportional to the total map activity is directly subtracted.

The dynamical equations of the pointer-map network are

$$\dot{M}_x = -M_x + [m_x - \beta \sum_{y=1}^N M_y + \alpha \sum_{k=1}^2 w_{xk} P_k]_+ \quad (2)$$

$$\dot{P} = -P + [p + \alpha \sum_{y=1}^N w_y M_y]_+ \quad (3)$$

where $[\cdot]_+ = \max(\cdot, 0)$ is the linear threshold or rectification nonlinearity, and the dot denotes temporal differentiation. α and β are positive parameters that determine excitatory and inhibitory connection strengths. The synaptic weights are symmetric between the map and the pointer, and are chosen to be excitatory and lie on a circle, as can occur in biology³⁷.

$$w_y = (w_{y1}, w_{y2}) = (\cos \delta_y, \sin \delta_y), \text{ where } \delta_y = \frac{\pi(y-1)}{2(N-1)} \quad (4)$$

Neurons are initialized with nonnegative activities (firing rates) P and M . The rectification ensures that the activities remain nonnegative at all times.

Equation 2 is written in component form and indexed over x to emphasize the ordered spatial arrangement of the map neurons. Equation 3 is written in vectorial form to emphasize that, geometrically, $P = (P_1, P_2)$ forms a planar vector that points onto the map (Fig. 2).

Pointer-map Lyapunov function. A Lyapunov function asserts that the ‘energy’ of the system continuously decreases and that the network converges to a stationary state for all initial conditions. Such functions have been demonstrated for saturating neural networks with symmetric weights ($w_{ij} = w_{ji}$)^{38–40}. This approach can also be applied to the pointer-map network with steady inputs. For equations (2) and (3), the Lyapunov function is given by

$$L(M, P) = Q(M, P) - \sum_x m_x M_x - p_1 P_1 - p_2 P_2 \quad (5)$$

and the quadratic function Q is

$$2Q = \sum_x M_x^2 + P_1^2 + P_2^2 + \beta \sum_{x,y} M_x M_y - 2\alpha \sum_x M_x (P_1 \cos \delta_x + P_2 \sin \delta_x) \quad (6)$$

For the time-derivative of L , we can show

$$\dot{L} \leq -\sum_x \dot{M}_x^2 - \dot{P}_1^2 - \dot{P}_2^2 \leq 0 \quad (7)$$

Hence, the continuously differentiable function L does not increase over time. More specifically, L is strictly decreasing and $\dot{L} = 0$ only at fixed points of equations (2) and (3). Unlike Lyapunov functions defined for saturating neurons, we must establish that L is bounded below. A sufficient condition for the lower bound is that the quadratic terms Q are positive everywhere except at the origin:

$$Q(M, P) > 0 \text{ for } \max(M_1, \dots, M_N, P_1, P_2) > 0 \quad (8)$$

It is possible to show⁴¹ that condition (8) is satisfied for the following choice of α :

$$\alpha \leq \sqrt{\frac{1}{N} + \beta} \quad (9)$$

This establishes that when α is chosen according to the ‘operational range’ (9) then for time-independent inputs m and p , every network trajectory converges to a fixed point. This operating point for α in (9) is safely below the true limit of stability of the network, which can be investigated by computer simulations.

Relation to recurrent maps. If the input to pointer neurons is positive, $p = p(\cos \theta, \sin \theta)$, where $p \geq 0$ and $0 \leq \theta \leq \pi/2$, then the pointer-map feedback loop can be reduced to regular local excitatory feedback only among map neurons. We assume that the pointer neurons have an initial state of direction γ_0 , given by $P(0) = P_0(\cos \gamma_0, \sin \gamma_0)$, $0 \leq \gamma_0 \leq \pi/2$. Because the pointer neurons receive purely excitatory input, there is no need for rectification. Hence, equation 3 can be integrated and the resulting expression for $P(t)$ can be placed into equation (2), which leads to:

$$\dot{M}_x = -M_x + [m_x - \beta \sum_y M_y + T_{1x} + T_{2x} + T_{3x}]_+ \quad (10)$$

Three terms, T_1 , T_2 and T_3 , describe inputs to the map and arise from its interaction with the pointer neurons:

$T_{1x} = e^{-t} P_0 \cos(\gamma_0 - \delta_x)$ is a preattentive input to the map arising from the initial setting of the pointer. This input is largest for the map neuron whose preferred stimulus location δ_x is closest to the initial pointing γ_0 , and decays to zero in time according to e^{-t} .

$T_{2x} = \alpha(1 - e^{-t})p \cos(\theta - \delta_x)$ is an attentional input to the map arising from the low-pass filtered input to the pointer. T_{2x} is initially zero and increases with time. Like T_{1x} , T_{2x} is a feedforward input to the map. Having a cosine profile, this input is maximal for the neuron in the map whose angle δ_x is closest to the angle θ of the pointer input.

$T_{3x} = \alpha^2 \sum_y \cos(\delta_x - \delta_y) \dot{M}_y(t)$ is the relevant feedback term of pointer maps. $\dot{M}_y(t) = \int_0^t ds e^{-(t-s)} \dot{M}_y(s)$ is the low-pass-filtered map activity. The pointer-map feedback loop reduces to delayed excitation of cosine profile on a recurrent map.

Note that the stationary states of pointer maps without attentional processing and recurrent maps with cosine interactions are exactly the same. Deviations from equation (10) occur for negative pointer input. In this case, input to the pointer neurons can reduce the excitatory connection strength between the map and the pointer by inactivating pointer neurons.

ACKNOWLEDGEMENTS

We thank Kevan Martin, Jochen Braun and David Feinstein for comments and discussions. We acknowledge the support of the Swiss National Science Foundation SPP Program.

RECEIVED 22 DECEMBER 1998; ACCEPTED 14 JUNE 1999

1. Douglas, R. J., Martin, K. A. & Witteridge, D. A canonical microcircuit for neocortex. *Neural Comput.* **1**, 480–488 (1989).
2. Douglas, R. J., Koch, C., Mahowald, M. A., Martin, K. A. & Suarez, H. Recurrent excitation in neocortical circuits. *Science* **269**, 981–985 (1995).
3. Douglas, R. J. & Martin, K. A. A functional microcircuit for cat visual cortex. *J. Physiol. (Lond.)* **440**, 735–769 (1991).
4. Nelson, S. B., Toth, L. J., Sheth, B. & Sur, M. Orientation selectivity of cortical neurons during intracellular blockade of inhibition. *Science* **265**, 774–777 (1994).
5. Ferster, D., Chung, S. & Wheat, H. Orientation selectivity of thalamic input to simple cells of cat visual cortex. *Nature* **380**, 249–252 (1996).
6. LeVay, S. & Gilbert, C. D. Laminar patterns of geniculocortical projection in the cat. *Brain Res.* **113**, 1–19 (1976).
7. Peters, A. & Payne, B. R. Numerical relationships between geniculocortical afferents and pyramidal cell modules in cat primary visual cortex. *Cereb. Cortex* **3**, 69–78 (1993).
8. Ahmed, B., Anderson, J. C., Douglas, R. J., Martin, K. A. & Nelson, C. Polynuclear innervation of spiny stellate neurons in cat visual cortex. *J. Comp. Neurol.* **341**, 39–49 (1994).
9. Somers, D. C., Nelson, S. B. & Sur, M. An emergent model of orientation selectivity in cat visual cortical simple cells. *J. Neurosci.* **15**, 5448–5465 (1995).
10. Ben-Yishai, R., Lev Bar-Or, R. & Sompolinsky, H. Theory of orientation tuning in visual cortex. *Proc. Natl. Acad. Sci. USA* **92**, 3844–3848 (1995).
11. Salinas, E. & Abbott, L. F. A model of multiplicative neural responses in parietal cortex. *Proc. Natl. Acad. Sci. USA* **93**, 11956–11961 (1996).
12. Chance, F. S., Nelson, S. B. & Abbott, L. F. Complex cells as cortically amplified simple cells. *Nat. Neurosci.* **2**, 277–282 (1999).
13. Pouget, A., Zhang, K., Deneve, S. & Latham, P. Statistically efficient estimation using population coding. *Neural Comput.* **10**, 373–401 (1998).
14. Moran, J. & Desimone, R. Selective attention gates visual processing in the extrastriate cortex. *Science* **229**, 782–784 (1985).
15. Luck, S. J., Chelazzi, L., Hillyard, S. A. & Desimone, R. Neural mechanisms of spatial selective attention in areas V1, V2, and V4 of macaque visual cortex. *J. Neurophysiol.* **77**, 24–42 (1997).
16. Colby, C. L., Duhamel, J.-R. & Goldberg, M. E. Visual, presaccadic, and cognitive activation of single neurons in monkey lateral intraparietal area. *J. Neurophysiol.* **76**, 2841–2852 (1996).
17. Treue, S. & Maunsell, J. H. Attentional modulation of visual motion processing in cortical areas MT and MST. *Nature* **382**, 539–541 (1996).
18. Fuster, J. M. Inferotemporal units in selective visual attention and short-term memory. *J. Neurophysiol.* **64**, 681–697 (1990).
19. Georgopoulos, A. P., Schwartz, A. W. & Kettner, R. E. Neuronal population coding of movement direction. *Science* **233**, 1416–1420 (1986).
20. Groh, J. M., Born, R. T. & Newsome, W. T. How is a sensory map read out? effects of microstimulation in visual area MT on saccades and smooth pursuit eye movements. *J. Neurosci.* **11**, 4312–4330 (1997).
21. Pouget, A. & Sejnowski, T. J. in *The Handbook of Brain Theory and Neural Networks* (ed., Arbib, M.) 335–338. (MIT Press, London, 1995).
22. Georgopoulos, A. P., Lurito, J., Petrides, M., Schwartz, A. B. & Massey, J. T. Mental rotation of the neuronal population vector. *Science* **243**, 234–236 (1989).
23. Snyder, L. H., Batista, A. P. & Andersen, R. A. Coding of intention in the posterior parietal cortex. *Nature* **386**, 167–170 (1997).
24. Andersen, R. A., Essick, G. E. & Siegel, R. M. Encoding of spatial location by posterior parietal neurons. *Science* **230**, 456–458 (1985).
25. Galletti, C. & Battaglini, P. P. Gaze-dependent visual neurons in area V3A of monkey prestriate cortex. *J. Neurosci.* **9**, 1112–1125 (1989).
26. Hansel, D. & Sompolinsky, H. in *Methods in Neuronal Modeling* 2nd edn., Ch. 13 (eds. Koch, C. & Segev, I.) (MIT Press Cambridge, Massachusetts, 1997).
27. Ballard, D. H., Hayhoe, M. M., Pook, P. K. & Rao, R. P. Deictic codes for the embodiment of cognition. *Behav. Brain Sci.* **20**, 723–757 (1997).
28. Hupé, J. M., James, A. C., Payne, B. R., Lomber, S. G., Girard, P. & Bullier, J. Cortical feedback improves discrimination between figure and background by V1, V2 and V3. *Nature* **394**, 784–787 (1998).
29. Motter, B. C. Focal attention produces spatially selective processing in visual cortical areas V1, V2, and V4 in the presence of competing stimuli. *J. Neurophysiol.* **70**, 909–919 (1993).
30. Reynolds, J. H., Chelazzi, L. & Desimone, R. Competitive mechanisms subserve attention in macaque areas V2 and V4. *J. Neurosci.* **19**, 1736–1753 (1999).
31. Luck, S. J. & Ford, M. A. On the role of selective attention in visual perception. *Proc. Natl. Acad. Sci. USA* **95**, 825–830 (1998).
32. Olshausen, B. A., Anderson, C. H. & Van Essen, D. C. A neurobiological model of visual attention and invariant pattern recognition based on dynamic routing of information. *J. Neurosci.* **13**, 4700–4719 (1995).
33. Salinas, E. & Abbott, L. F. Invariant visual perception from attentional gainfields. *J. Neurophysiol.* **77**, 3267–3272 (1997).
34. Andersen, R. A., Bracewell, M. R., Barash, S., Gnadt, J. W. & Fogassi, L. Eye position effects on visual, memory and saccade-related activity in areas LIP and 7a of macaque. *J. Neurosci.* **10**, 1176–1196 (1990).
35. Gottlieb, J. P., Kusunoki, M. & Goldberg, M. E. The representation of visual salience in monkey parietal cortex. *Nature* **391**, 481–484 (1998).
36. Kernighan, B. W. & Ritchie, D. M. *The C Programming Language* (Prentice Hall, Englewood Cliffs, New Jersey, 1978).
37. Lewis, J. E. & Kristan, W. B. Jr. A neuronal network for computing population vectors in the leech. *Nature* **391**, 76–79 (1998).
38. Hopfield, J. J. Neurons with graded response have collective properties like those of two-state neurons. *Proc. Natl. Acad. Sci. USA* **81**, 3088–3092 (1984).
39. Cohen, M. A. & Grossberg, S. Absolute stability of global pattern formation and parallel memory storage by competitive neural networks. *IEEE Trans. Sys. Man Cybern.* **13**, 288–307 (1983).
40. Grossberg, S. Nonlinear neural networks: principles, mechanisms, and architectures. *Neural Net.* **1**, 17–61 (1988).
41. Hahnloser, R. H. R. Computation in Recurrent Networks of Linear Threshold Neurons: Theory, Simulation and Hardware Implementation. Thesis, Eidgenössische Technische Hochschule Zürich (1998).